A CONTROL SYSTEM ARCHITECTURE FOR INTELLIGENT MACHINE SYSTEMS

James S. Albus
Robot Systems Division
National Bureau of Standards

I. Introduction

An intelligent machine system is an autonomous or semi-autonomous system that can do useful work in an uncertain environment. An intelligent machine system consists of the eight basic elements shown in Figure 1.

1. ACTUATORS -- The output of an intelligent machine system is produced by a set of actuators and servo controls which position, move, and exert forces on arms, hands, tools, vehicles, and sensor platforms. An intelligent machine system may have tens, hundreds, or even thousands of actuators, all coordinated in order to accomplish a goal or task.

2. SENSORS -- Input is through a set of sensors such as vision, tactile, force, torque, position, distance, vibration, acoustic, and temperature measuring devices that monitor both the state of the external world and the internal state of the intelligent machine system.

3. TASK DECOMPOSITION -- Behavior is generated by a system which decomposes high level tasks into low level actions. An intelligent task decomposition system has the ability to formulate or select effective plans of action and set priorities based on values such as cost, risk, and utility. Plans are strings of intended actions selected from a wide range of potential action sequences. An intelligent task decomposition system is able to plan and execute effective actions in the presence of uncertain, incomplete, and sometimes incorrect information. An intelligent system continuously monitors the execution of plans, and replans when the situation requires. It also has the ability to reason about geometry, dynamics, and space-time.

Commerical equipment is identified in this paper in order to adequately describe the systems that were developed. In no case does such identification imply recommendation by the National Bureau of Standards, nor does it imply that this equipment was necessarily the best available for the purpose.

The example task decomposition system described in this paper consists of a hierarchy of planners that simultaneously generate and coordinate plans at each level, with different planning horizons and different degrees of detail at each level.

4. GOAL SELECTION -- Goals are defined as situations or events to be achieved which have value to the system. At the highest level of control, an intelligent machine is a system which selects and prioritizes goals so as to maximize the benefit and minimize the cost and risk for the overall system.

5. VALUES -- In order to select goals and assign priorities, an intelligent machine system must have a set of values which enable intelligent machine systems to evaluate situations and events. These values enable the control system to choose between goals, to assign resources to tasks, and to generate priorities that maximize payoff and minimize cost and risk. The evaluation of both planned and observed situations and events depends on the values assigned to resources such as time and material, as well as on the priorities assigned to cost, risk, payoff, and survival.

6. WORLD MODEL -- An intelligent machine system must have an internal model of the world which contains knowledge about:

a) both generic and specific objects, groups, and assemblies in the world. This knowledge includes position, velocity, geometry, topology, material properties, surface characteristics, dimensional tolerance, and kinematic and dynamic properties. About intelligent objects, the world model may also contain knowledge of capabilities, typical behaviors, and probable intentions.

b) the world itself including both specific information such as temperature, wind velocity, and illumination, as well as general information such as the laws of physics, chemistry, optics,

c) 2-D and 3-D maps of how space is filled, both from a world coordinate frame and from each individual machine's viewpoint,

d) measures of probability, confidence, and believability for data in the world model.

e) The world model must also have means for updating and maintaining consistency of its databases, and for making predictions as to what sensory data is expected, and for simulating what results could be anticipated from hypothesized plans of action.

7. SENSORY PROCESSING -- An intelligent machine system must have a set of sensory signal processing algorithms that detect events, and recognize features, objects, and relationships in the

world. An intelligent machine system must integrate both temporally and spatially over signals derived from a wide variety of sensors so as to determine distance, shape, orientation, surface characteristics, and material properties of objects and regions of space. Sensory processing may include recognition of speech, and interpretation of visual images, language, and music.

8. SYSTEM INTEGRATION -- In order to connect the above systems together, there must be a computing architecture, distributed databases, communications, and operator interface. These should include systems for programming, debugging, generating process plans, and entering geometric data describing objects and terrain data of regions describing regions of the world.

As defined here, the superstructure of intelligent machine systems is generic and may be applied to a wide variety of application areas, including manufacturing, construction, undersea, military, outer space, and domestic service robots. Each specific intelligent machine system will have application specific data in its world model, and domain specific algorithms in its control and sensory processing systems.

What is common to all intelligent machine systems is the ability to process sensory information, to build and use world models, to decompose tasks, to plan, schedule, and execute actions, and to evaluate the results of those actions. The world model of an intelligent machine system will contain much apriori knowledge such as maps of places and the geometry of objects. The sensors and sensor processing systems allow this apriori information to be verified, updated, modified, or deleted based on observations.

An intelligent machine system may include robots, machine tools, materials transport systems, assembly machines, or inspection systems which move, position, cut, fasten, push, and twist tools and objects. It will typically contain some mobility mechanism or vehicle, and some type of manipulator arms, hands, and tools. It will necessarily have many sensors, and a real-time, sensory-interactive, goal-directed control system.

II. An Experimental Intelligent Machine System Multiple Autonomous Underwater Vehicles

The control system architecture being developed for the Multiple Autonomous Underwater Vehicle (MAUV) project is an example of an intelligent machine system. The specific research topics being addressed by the MAUV project are: hierarchical distributed control, real-time planning, world modeling, value-driven reasoning, intelligent sensing and communication, and cooperative problem solving by two intelligent vehicles in a natural and potentially hostile environment.

Demonstration scenarios have been designed to study, and attempt to mimic, aggression, predation, exploration, escape, communication, and cooperation. Such behavioral activities are common in all intelligent creatures in nature, and are particularly refined in humans. It can be argued that relative success or failure in these behaviors is what provides the natural selection pressure to drive the evolution of intelligence.

Among the behavioral scenarios being developed are:

a) One vehicle searches an area while the other relays messages about what has been found.

b) One vehicle illuminates a target while another takes pictures.

c) One vehicle hunts for prey while the other lies in wait.

d) One vehicle attracts the prey's attention, while the other closes in for the kill.

e) When in danger, one vehicle draws attention to itself, while the other escapes with valuable information.

Value driven logic [1] allows each vehicle to weigh the value of its own survival against the success of the mission.

The study of scenarios involving two vehicles permits research on communication as a goal directed activity. For example, what information should be transmitted? for what purpose? and when? When is the value of a piece of information of sufficient value to incur the risk to survival of revealing one's presence by transmitting a message? What are communication strategies which balance the risk against the benefits?"

Issues of command and control can also be studied. For example, how should control systems be structured so that both vehicles can behave equally intelligently when they are apart, but one vehicle is recognized as the leader of the group when they are together? How do they share knowledge acquired by only one? What do they do if they cannot agree on a strategy?

The autonomous undersea vehicles selected for the MAUV project were designed and constructed by the University of New Hampshire Marine Systems Engineering Laboratory. They are based on the EAVE-EAST design [2]. The MAUV control system architecture is based on the NASA/NBS standard reference model control system architecture (NASREM) [3]. The MAUV project is funded by the DARPA Naval Technology Office. Figure 2 is a picture, and Figure 3 a diagram, of a MAUV vehicle. The vehicle is gravity stabilized in pitch and roll, with thrusters that allow it to be controlled in x, y, z, and yaw. It is a battery powered with the batteries stored in cylindrical tanks at the bottom of the vehicle, and flotation tank on the upper part of the vehicle. Each vehicle carries three acoustic navigation transponders which allow it to measure the range and bearing to navigation bouys placed in the water. The vehicle has a compass, pressure and temperature sensors, and bottom and surface sounders. In front, it has an obstacle avoidance sonar consisting of five narrow beam acoustic transmitter-receivers. These are arranged such that the center sonar beam points straight ahead, two point ten degrees to the right and left, and two point ten degrees up and down from the center beam. Each vehicle also carries an radio frequency communications system and in the future will also have an acoustic communications system. Additional sensors planned for the future include scanning imaging sonar, vision, and various structured lighting techniques.

III. THE MAUV CONTROL SYSTEM ARCHITECTURE

The MAUV control system architecture incorporates a number of concepts developed in previous and on-going robotics research programs, including the NASA telerobotics program [3], the DARPA Autonomous Land Vehicle [4], the Air Force/DARPA Intelligent Task Automation program [5], the supervisory control concepts pioneered by Sheridan at MIT [6], and the hierarchical control system developed for the Automated Manufacturing Research Facility at the National Bureau of Standards [7-9]. The MAUV architecture integrates many artificial intelligence concepts such as goal decomposition, hierarchical real-time planning, model driven image analysis [10-14], blackboards [14], and expert systems into a systems framework with modern control concepts such as multivarient state space control, reference model adaptive control, dynamic optimization, and learning systems [15-18]. The MAUV architecture framework also readily accommodates concepts from operations research, differential games, utility theory, and value driven reasoning [1, 19-20]

A block diagram of the NBS MAUV control system architecture is shown in Figure 4. In the MAUV control system architecture the task decomposition modules perform real-time planning and task monitoring functions, and decompose task goals both spatially and temporally, as shown in Figure 5. The sensory processing modules filter, correlate, detect, and integrate sensory information over both space and time so as to recognize and measure patterns, features, objects, events, and relationships in the external world. The world modeling modules answer queries, make predictions, and compute evaluation functions on the state space defined by the information stored in global memory, as shown in Figure 6. Global memory is a database which contains the system's best estimate of the state of the external world. The world modeling modules keep the global memory database current and consistent.

Task Decomposition - H modules
(Plan, Execute)

The task decomposition hierarchy consists of H modules which plan and execute the decomposition of high level goals into low level actions. Task decomposition involves both a spatial decomposition (into concurrent actions by different subsystems), and a temporal decomposition (into sequential actions along the time line).

Each H module at each level consists of three sublevels as shown in Figure 4:

1) a planner manager PM

2) a set of planners PL(i) and

3) a set of executors EX(i).

These three sublevels decompose the input task into both spatially and temporally distinct subtasks as shown in Figure 4.

For each level:

Planner Manager

The planner manager PM is responsible for partitioning the task command into i spatially or logically distinct jobs to be performed by i physically subsystems. Each subsystem has a planner/executor mechanisms. At the upper levels the job assignment module may also assign physical resources to the subsystems for each job.

Planners

For each job assigned to a subsystem, there exists a planner PL(i) and a executor EX(i). Each planner is responsible for decomposing its job assignment into a temporal sequence of planned subtasks to be executed by the respective executor.

The planning horizon is defined as the period into the future over which a plan is prepared. Each level of the hierarchy has a planning horizon of approximately two input task time durations. Replanning is done at cyclic intervals, as well as whenever emergency conditions arise. The cyclic replanning interval is about an order of magnitude less than the planning horizon (or about equal to the expected output subtask time duration). Emergency replanning begins immediately upon the detection of emergency conditions.

Executors

Each executor EX(i) is responsible for successfully executing the plan prepared by its respective planner PL(i). If all the subtasks in the plan are successfully executed, then the goal of the original task will be achieved. The executor operates by selecting the first subtask from the current queue of planned subtasks and outputting a subcommand to the appropriate subordinate H module at time t. The EX(i) module monitors its feedback input in order to servo its output to the desired subtask activity.

The feedback also carries subgoal event information and a time of day clock for coordination of output between executors at the same level. When the executor detects a subgoal event, it steps to the next subtask in the plan.

World Modeling - M modules
(Remember, Estimate, Predict, Evaluate)

The world modeling leg of the hierarchy consists of M modules which model (i.e. remember, estimate, predict) and evaluate the state of the world. The world model is the system's best estimate and evaluation of the history, current state, and possible future states of the world, including the states of the system being controlled. The world model includes both the M modules and a knowledge base stored in global memory where state variables, maps, lists of objects and events, and attributes of objects and events are maintained. By this definition, the world model corresponds to what is widely known in the artificial intelligence community as a "blackboard" [14].

As shown in Figure 6, the M modules at various levels:

a) Maintain the global memory knowledge base, keeping it current. The M modules update the knowledge base based on correlations and differences between model predictions and sensory observations.

b) Provide predictions of expected sensory input to the corresponding G modules, based on the state of the task and estimates of the external world.

c) Answer "What is?" questions asked by the planners and executors in the corresponding level H modules. The task executor requests information about the state of the world, and uses the answers to monitor and servo the task, and/or to branch on conditions to subtasks that accomplish the task goal.

d) Answer "What if?" questions asked by the planners in the corresponding level H modules. M modules predict the results of hypothesized actions.

e) Evaluate the current situation and potential future consequences of hypothesized actions by applying evaluation functions to current states and to future states expected to result from hypothesized actions. The evaluation functions include a set of values assigned to events such as vehicle survival, subtask completion, and information gathered by the vehicles. They also include a set of priorities assigned to each of these values.

The evaluation functions enable value driven decision logic [20] at several different hierarchical levels. Working together, the world model predictors, evaluation functions, and the planners are able to search the space of possible futures, and choose the sequence of planned actions that produce the best evaluation. The executors are also able to apply value driven logic to moment by moment behavioral decisions.

Global Memory

Global memory is the database in which knowledge is stored about the state of the world, including the internal state of the control system.

Contents of Global Memory

The knowledge in the global memory consists of:

a) Maps which describe the spatial occupancy of the world. A map is a spatially indexed database showing the relative position of objects and regions. At different levels the maps have different resolution. Resolution increases at each successively lower level, while area covered by the map increases at each successively higher level. The maps at different levels thus represent a pyramid structure. Maps may also contain a number of overlays. These overlays may indicate values such as utility, cost, risk, etc. to be used for planning.

b) Lists of all known objects, features, regions, relationships, and events are listed in the global memory database indexed by name, along with frames containing their attributes. The object and feature frames contain information such as position, velocity, orientation, shape, dimensions, reflectance, color, mass, and other information of interest. Event frames contain information such as start and end time, duration, type, cost, payoff, etc. Recognized objects and events may also have associated with them confidence levels, and degrees of believability and dimensional certainty.

At different levels, object frames have different levels of detail and spatial resolution, and event frames have different levels of temporal resolution.

c) State variables which are the system's best estimate of the state of the world, including both the external environment and the internal state of the H, M, and G modules. Data in global memory is available to all modules at all levels of the control system.

Implemention of Global Memory

Global memory in the MAUV architecture is not located in a single physical database, but is distributed over several computers, memory boards, and mass storage devices on a VME bus. Global memory is, in fact, distributed over more than one vehicle. Variables in global memory are globally defined, i.e., they may be accessed (read or written) by name from local processes running at any level. Of course, the time required to access a global variable is not the same for all processes. For example, in order for a global variable in vehicle-A to be read or updated by a process in vehicle-B, the two vehicles may have to rendezvous and communicate world model updates. This may take many minutes or hours.

Sensory Processing - G modules
(Filter, Integrate, Detect, Measure)

The sensory processing leg of the MAUV control hierarchy consists of G modules which recognize patterns, detect events, and filter and integrate sensory information over space and time. The G modules are similar to the H modules in that they also consist of three sublevels which:

1) compare sensor observations with world model predictions

2) integrate correlation and difference over time

3) integrate correlation and difference over space

These spatial and temporal integrations fuse sensory information from multiple sources over extended time intervals. Newly detected or recognized events, objects, and relationships are entered by the M modules into the world model global memory database, and objects or relationships perceived to no longer exist are removed. The G modules also contain functions which can compute confidence factors and probabilities of recognized events, and statistical estimates of stochastic state variable values.

VI. Functional Levels in the MAUV Hierarchy

The MAUV control system hierarchy is structured into six layers, as shown in Figure 3 and 6. At each layer, a different fundamental transformation is performed.

Level 1 -- Coordinate Transform/Servo
transforms coordinates from a vehicle coordinate frame into actuator coordinates. This level also servos actuator power.

Level 2 -- Dynamic (Primitive)

works in vehicle or world coordinates. It computes inertial dynamics, and generates smooth trajectory positions, velocities, accelerations for efficient vehicle maneuvers.

Level 3 -- Elementary Move (E-Move)

works in both symbolic and geometric space. It decomposes elementary move commands (E-moves) into strings of intermediate poses, or dynamic (primitive) level commands.

As shown in Figure 7, each MAUV vehicle consists of three subsystems: pilot, communications, and sonar. E-moves are defined for each vehicle subsystem.

A pilot E-Move can be defined as a smooth coordinated motion of the vehicle designed to achieve some position, orientation, or "key-frame pose" in state-space, or space-time. The level 3 pilot planner computes clearance with obstacles sensed by on-board sensors and generates strings of intermediate poses that define motion pathways between key-frame poses.

A communications E-move is a message. The level 3 communications planner encodes messages into strings of symbols, adds redundancy for error detection and correction, and formats the symbols for transmission.

A sonar E-Move may be defined as a temporal pattern of sonar pings or a scanning pattern for a passive listening beam designed to obtain a specific type of information about a specific target. The level 3 sonar planner decomposes sonar E-Moves into patterns of sonar pings and scanning beam dwell times.

Level 4 -- Vehicle Task

Level 4 works in object/task space. It decomposes vehicle commands, defined in terms of tasks to be performed by a single AUV on a target object, into sequences of E-moves, defined in terms of vehicle subsystem actions on aspects, or features, of an object.

The level 4 planner manager decomposes vehicle tasks into work elements to be performed by the various vehicle subsystems. It also coordinates, synchronizes, and resolves conflicts between vehicle subsystem plans.

The level 4 planners schedule sequences of E-Moves for the pilot, the communications, and the sonar subsystems.

The level 4 pilot planner checks the world model map to assure that there exists at least one pathway between keyframe poses. From the map, it estimates the cost, risk, and benefit of various routes and chooses a path that maximizes some cost-benefit evaluation function.

The level 4 communications planner schedules the messages to be sent. It computes the value of each message, its urgency, the risk of breaking communications silence, the power needed to make the message heard, and decides if and when to send the message.

The level 4 sonar planner analyzes the nature of the target, plans scanning patterns for passive or active beams, estimates the value of taking an active sonar sounding, and compares that against the risk of breaking silence.

Level 5 -- Group

Level 5 decomposes group tasks into vehicle tasks. Group tasks define actions to be performed on groups of objects by groups of autonomous vehicles. Level 5 decomposes these into sequences of tasks for individual vehicles to perform on individual objects. The level 5 route planners use a Group level world model map to compute vehicle transit times. Level 5 also estimates costs, risks, and benefits of various vehicle task sequences. It schedules the actions of each AUV to coordinate with the other AUV in the group so as to maximize the effectiveness of the MAUV group.

Level 6 -- Mission

Level 6 decomposes a commanded MAUV mission into a sequence of group tasks and assigns priorities and values to various group tasks and mission subtasks. Missions are typically specified by a list of mission objectives, priorities, requirements, and time line constraints. The level 6 planning manager assigns mission objectives to MAUV groups. The level 6 planner generates requirements for resources such as fuel, and time, develops a schedule, and sets priorities for each respective group assignment. It schedules the activities of the group so as to maximize the effectiveness of the total mission.

V. Real Time Planning

One of the unique features of the MAUV hierarchical computing architecture is its multiple levels of planners. These provide a unique ability to deal with the real-time aspects of planning and task execution. Because planning is done at each level, plans at any one level typically consist of less than ten steps, and hence can be derived relatively quickly.

Planning is done top-down. The highest level plan covers the entire backlog of tasks to be accomplished by the end of the mission. At each lower level, plans are formulated (or selected) in real-time to accomplish the next step (or two) in the plan of the level immediately above.

Figure 8 shows an example of three levels of hierarchical planning activity. The activity represented at the highest level is input to the top level H module as a task command. This task is decomposed by the job assignment manager and three planners of the top H module into three concurrent plans consisting of four activity-event pairs each. The first executor of the top level H module outputs the current subtask command in its plan to a second level H module.

At the second level, the input task command is further decomposed by the H module into three concurrent plans, each consisting of four subtasks each. At the third level, the H module further decomposes its input task into three plans of four subtasks.

At each level, the final subgoal events in the plans correspond to the goal of the input task. At each successively lower level, the planning horizon becomes shorter, and the subtask activities become more detailed and fine structured.

The MAUV control system always has a hierarchy of plans in place. The MAUV vehicles begin each mission with a mission plan, with a planning horizon to the end of the mission. The group level always has a plan for how to accomplish the next two steps in the mission plan, and so on down the hierarchy. If the mission goes as planned, each level of the control system will always be able to anticipate the next subtask, and there is no need to pause to replan. However, if unexpected events cause a plan at some level of the hierarchy to become obsolete, a new plan must be generated. If a new plan cannot be developed in time, the real-time executor will be without a plan. A condition in which one or more levels has no plan available for execution can be described as a state of confusion. The executor must then execute some preplanned emergency action until the emergency planner can generate a new plan.

VI. Implementation

The MAUV control architecture is being implemented on the computing systems shown in Figure 9. In each vehicle, a VME bus supports high bandwidth communication between sensory processing, world modeling, task planning, and task execution modules at each level of the hierarchy. The set of computing modules is partitioned between three separate single board computers so as to maximize the use of parallel computation. A two megabyte common memory board is used for communication between processes, and a 800 megabyte optical disk will be used for mass storage. pSOS and pRISM from Ironics is being used as the real-time multi-processor operating system.

Also shown in Figure 9 is the MAUV software development and simulation environment. A variety of software development tools, such as Sun workstations, a VAX 11/785, a micro-VAX, and Iris graphics systems, PC's and Duals are being tied into the development environment for code development and simulation. Translators and cross compilers are being provided so that software developed in this environment can be downloaded into the 68020 target hardware for real-time execution.

The MAUV computing architecture can accommodate many additional computers as the complexity of the tasks assigned to the vehicles is increased. In the future, as vision and sonar arrays are added to the vehicles, the MAUV computer systems will incorporate special purpose computing elements, such as pipeline image processors and vector accelerators. These will physically interface with the control system through the VME bus.

VII. Conclusion

The MAUV project is an experimental intelligent machine system. It incorporates each of the eight elements described in section I. The the current implementation, however, is very preliminary and incomplete. Much additional work remains to increase the sophistication of the planners and executors at all levels, and to improve the world model representation of geometry, topology, and spatial relationships in maps and object lists. Many new sensors and sensory processing algorithms need to be added.

The architecture for intelligent machine systems described here provides the framework for integrating these additional capabilities. Experience to date suggests that this architecture provides a framework for the rapid development of intelligent machine systems.

ACKNOWLEDGEMENT:
This project is funded by the Defense Advanced Research Projects Agency's (DARPA) Naval Technology Office. The persons working on the MAUV from the National Bureau of Standards are: A.Abrishamiam, M. Ali, S. Chang, K. Goodwin, M. Herman, T. Hong, H. Huang, M. Juberts, J. Lee, P. Mansbach, J. Michaloski, D. Orser, D. Oskard, R. Quintero, M. Roche, S. Swetz, T. Tsai, B. Warsaw, and T. Wheatley. The persons working on this project from the University of New Hampshire: D. Blidberg, S. Chappell, R. Haywood, J. Jalbert, M. Shevenell, D. Stamulis, B. Welsh, and other support staff and students. Contractors involved in this project include: Decision-Science Applications, Inc., Martin Marietta-Baltimore, and Robotic Technology, Inc.

VIII. REFERENCES

[1] George Pugh, G.L. Lucas, "Application of Value-Driven Decision Theory to the Control and Coordination of Advanced Tactical Air Control Systems", Decision-Science Applications,Inc., Report No. 218, April 1980.

[2] D. R. Blidberg, "Guidance Control Architecture for the Eave Vehicle", IEEE Journal Oceanic Engineering, October 1986.

[3] J. S. Albus, H. G. McCain, and R. Lumia, "NASA/NBS Standard Reference Model for Telerobot Control System Architecture (NASREM)", NBS Technical Report

[4] J. Lowrie, et.al. "Autonomous Land Vehicle" Annual Report, ETL-0413, Martin Marietta Denver Aerospace, July 1986.

[5] John Graham, "Intelligent Task Automation (Phase II) Sixth Quarter Quarterly Report", Martin Marietta Denver Aerospace, March, 1987.

[6] T.B. Sheridan "Supervisory Control of Remote Manipulators, Vehicles and Dynamic Processes". In Rouse, W.B. (Ed.) Advances in Man-Machine Systems Research, Vol. 1, NY JAI Press, 49-137, 1984.

[7] J.A. Simpson, R.J. Hocken, J.S. Albus, "The Automated Manufacturing Research Facility of the National Bureau of Standards," Journal of Manufacturing System, Vol. 1, No. 1, 1983, p. 17.

[8] C. McLean, M. Mitchell, E. Barkmeyer, "A Computer Architecture for Small Batch Manufacturing," IEEE Spectrum, May, 1983, p. 59.

[9] J. Albus, C. McLean, A. Barbera, M. Fitzgerald, " An Architecture for Real-Time Sensory-Interactive Control of Robots in a Manufacturing Environment," 4th IFAC/IFIP Symposium on Information Control Problems in Manufacturing Technology, Gaithersburg, Oct., 1982.

[10] M.O. Shneier, E.W. Kent, P. Mansbach, "Representing Workspace and Model Knowledge for a Robot with Mobile Sensors," Proc. 7th Int. Conf. Pattern Recognition, 1984, p. 199.

[11] W.A. Perkins, "A Model Based Vision System for Industrial Parts," IEEE Trans. on Computers, Vol. C-27, 1978, p. 126.

[12] G.L. Gleason, G.J. Agin, "A Modular Vision System for Sensor-controlled Manipulation and Inspection," Proc. 9th Int. Symposium on Industrial Robots, 1979, p. 57.

[13] R.C. Bolles, P. Horaud, M.J. Hannah, "3DPO: Three Dimensional Parts Orientation System," Proc. of The Int. Joint Conf. on Artificial Intelligence, August 1983, p. 1116.

[14] T.F. Knoll and R.C. Jain, "Recognizing Partially Visible Objects Using Feature Indexed Hypotheses," Proc. IEEE Conf. on Robotics and Automation, San Francisco, 1986, p. 925.

[15] A. Barr, E. Feigenbaum, The Handbook of Artificial Intelligence, (Los Altos, William Kaufman, 1981).

[16] M. Brady, et.al., ed. Robot Motion: Planning and Control, (Cambridge, MIT Press, 1982).

[17] D.E. Whitney, "Resolved Motion Rate Control of Manipulators and Human Prostheses," IEEE Trans. Man-Machine Systems MMS-10, 1969, p. 47.

[18] R.P. Paul, Robot Manipulators: Mathematics, Programming, and Control, (Cambridge, MIT Press, 1981.)

[19] J.C. Crowley, "Navigation for an Intelligent Mobile Robot," IEEE Journal of Robotics and Automation, Vol. RA-1, No. 1, 1985, p. 31.

[20] George E. Pugh, D.F. Noble "An Information Fusion System for Wargaming and Information Warfare Applications", Decision-Science Applications, Inc., Report No. 314, May 1981

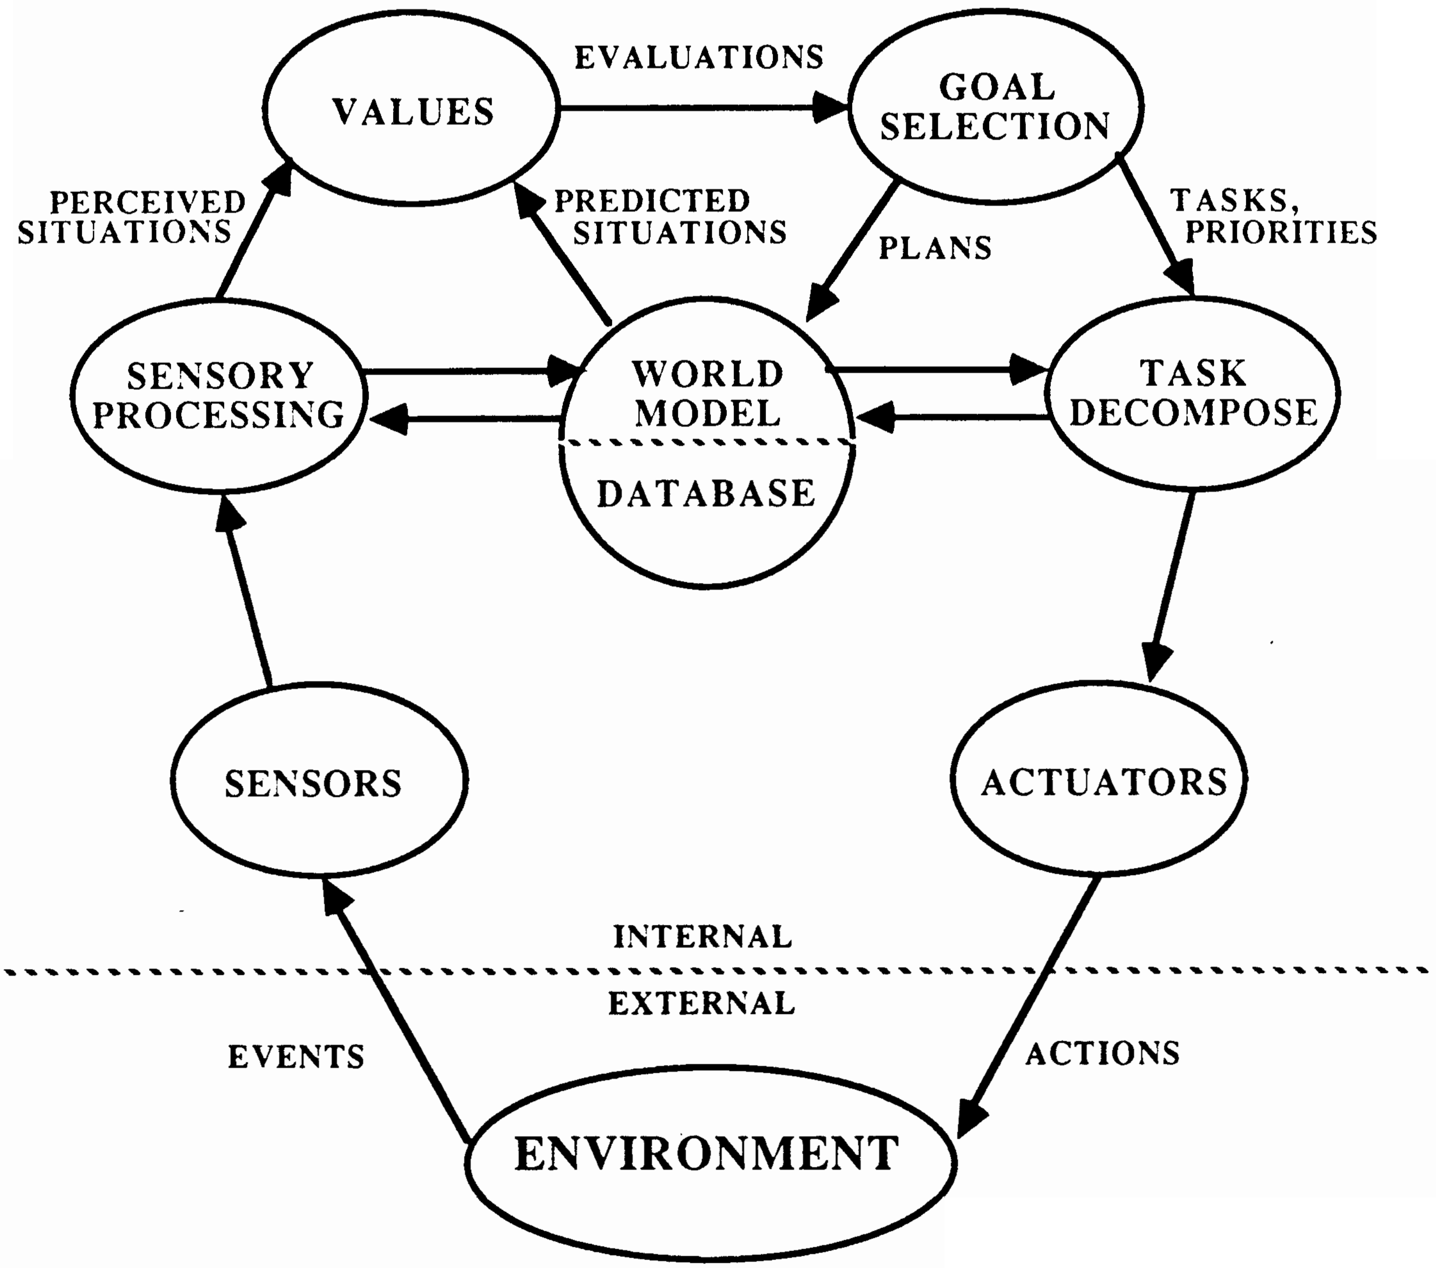

FIGURE 1: AN INTELLIGENT MACHINE SYSTEM

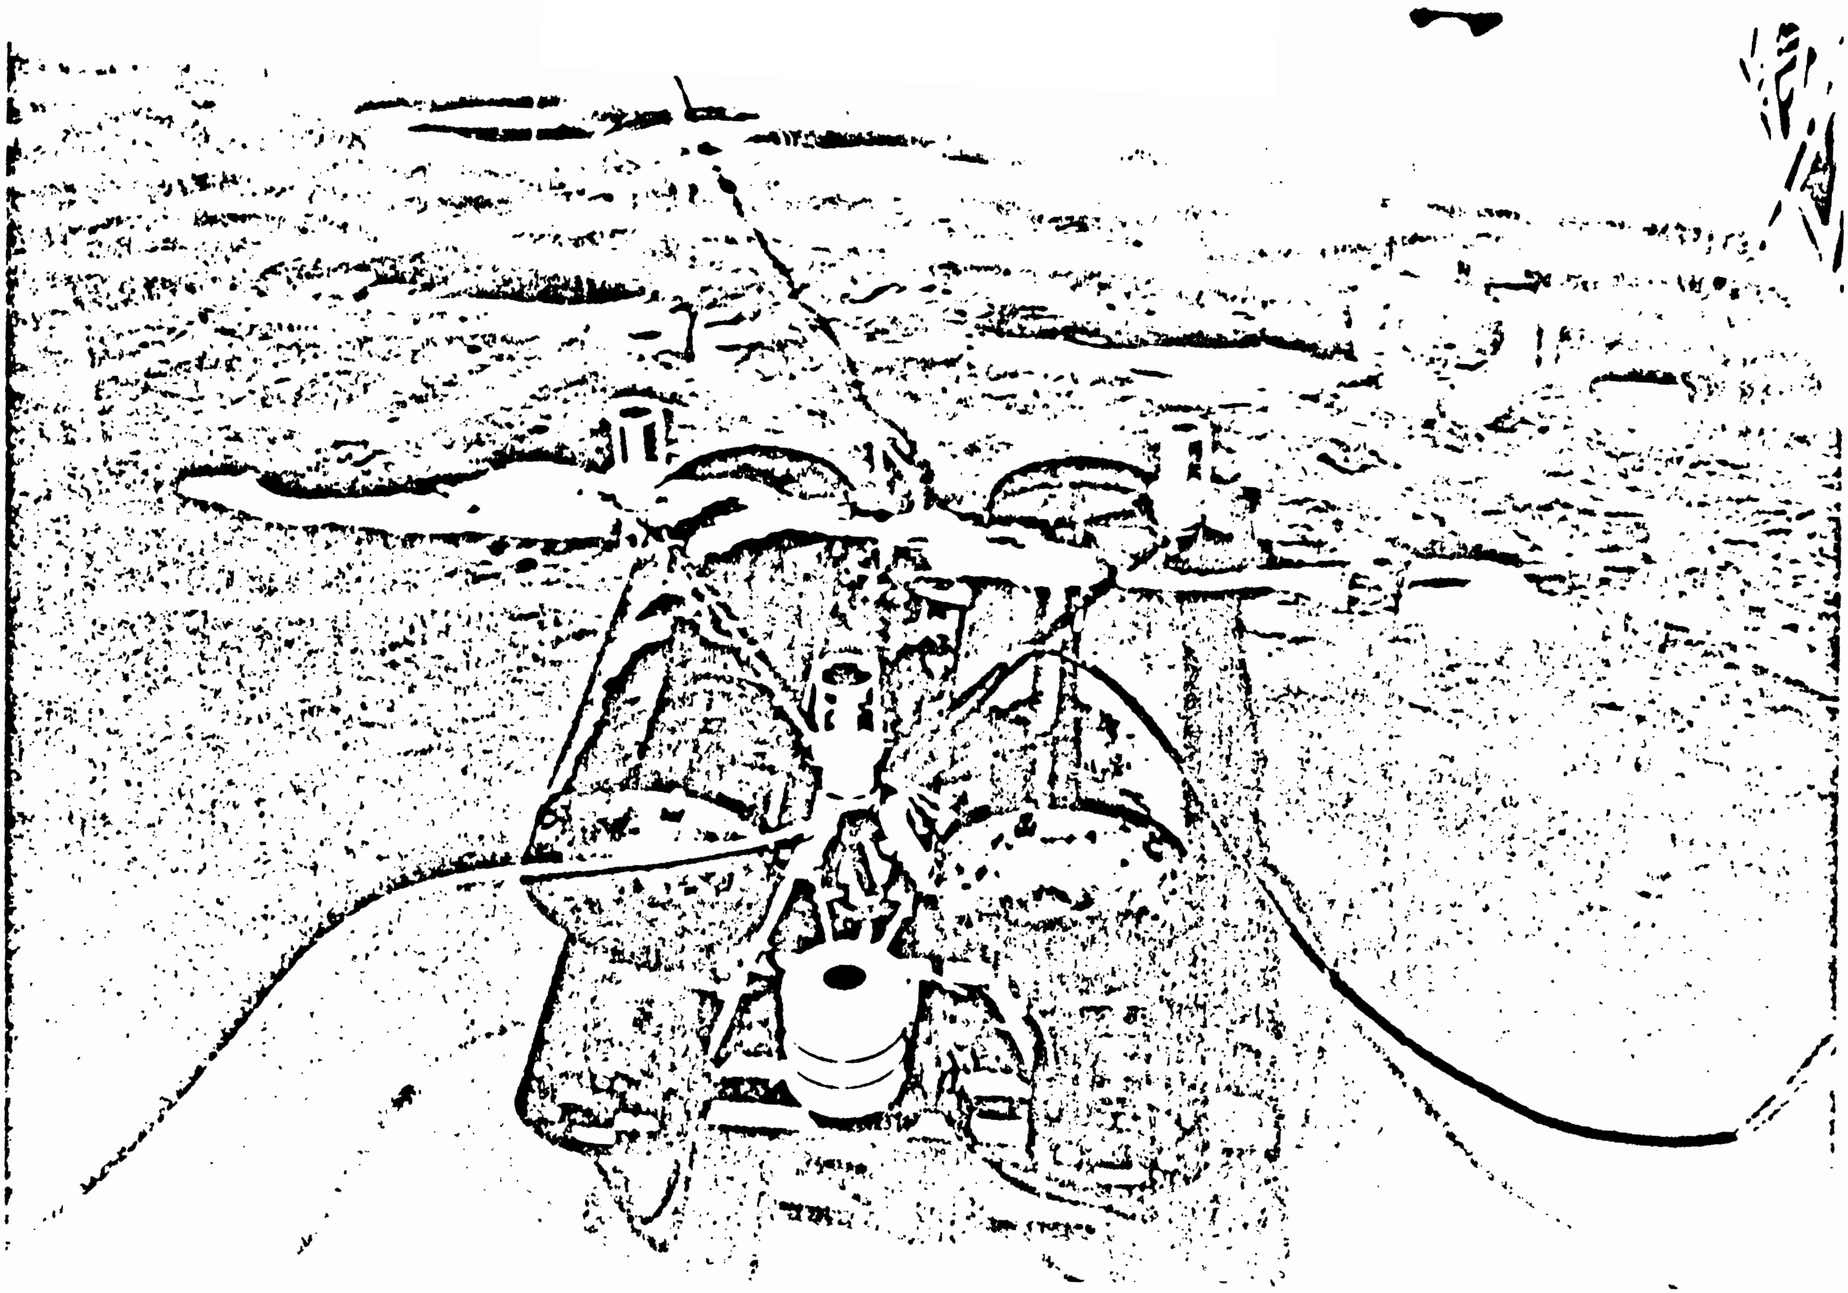

FIGURE 2: One of the MAUV Vehicles undergoing tests in a tank. Ropes and tethers will be removed during operational manuevers.

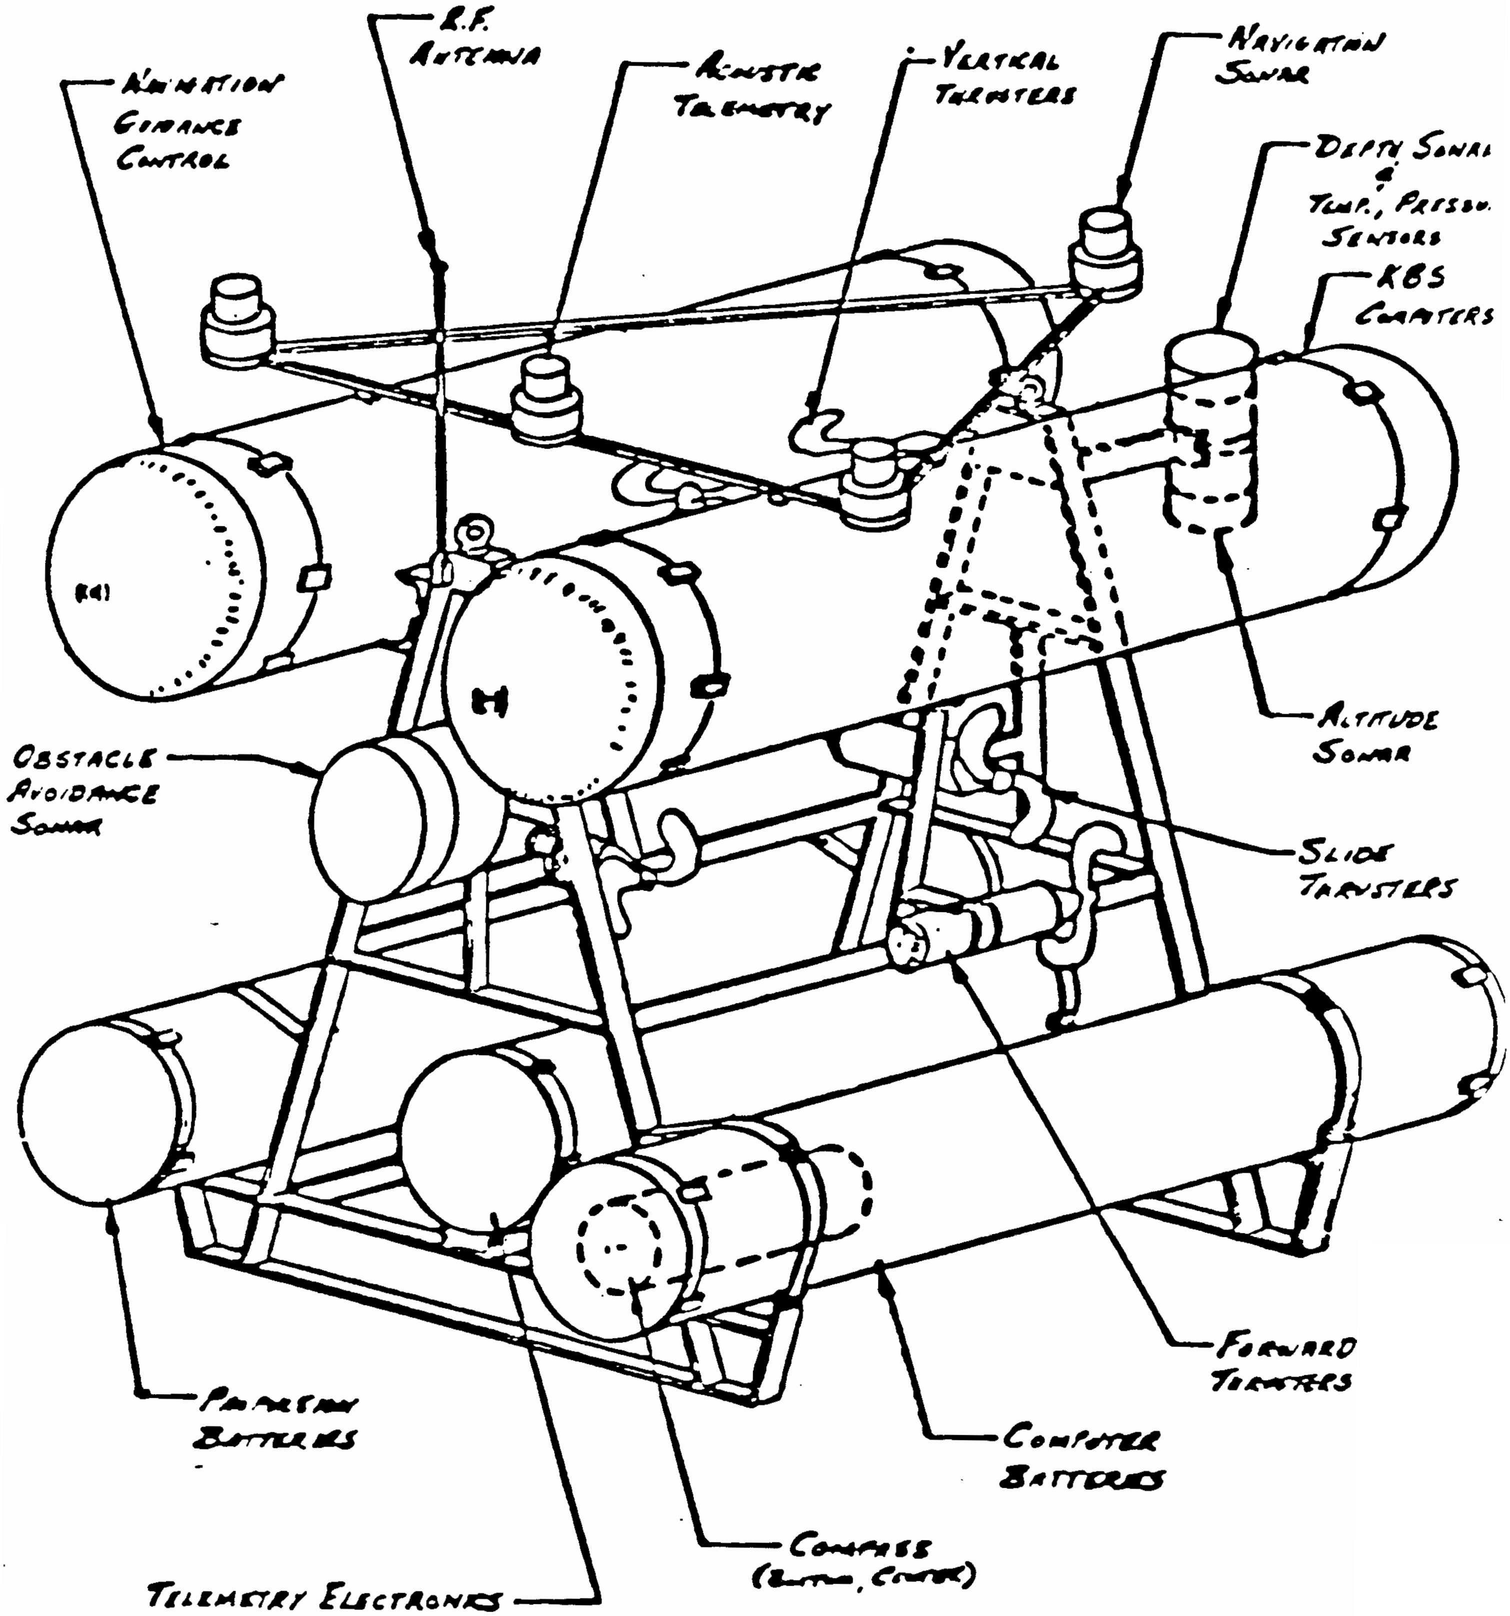

FIGURE 3: A diagram of a MAUV Vehicle.

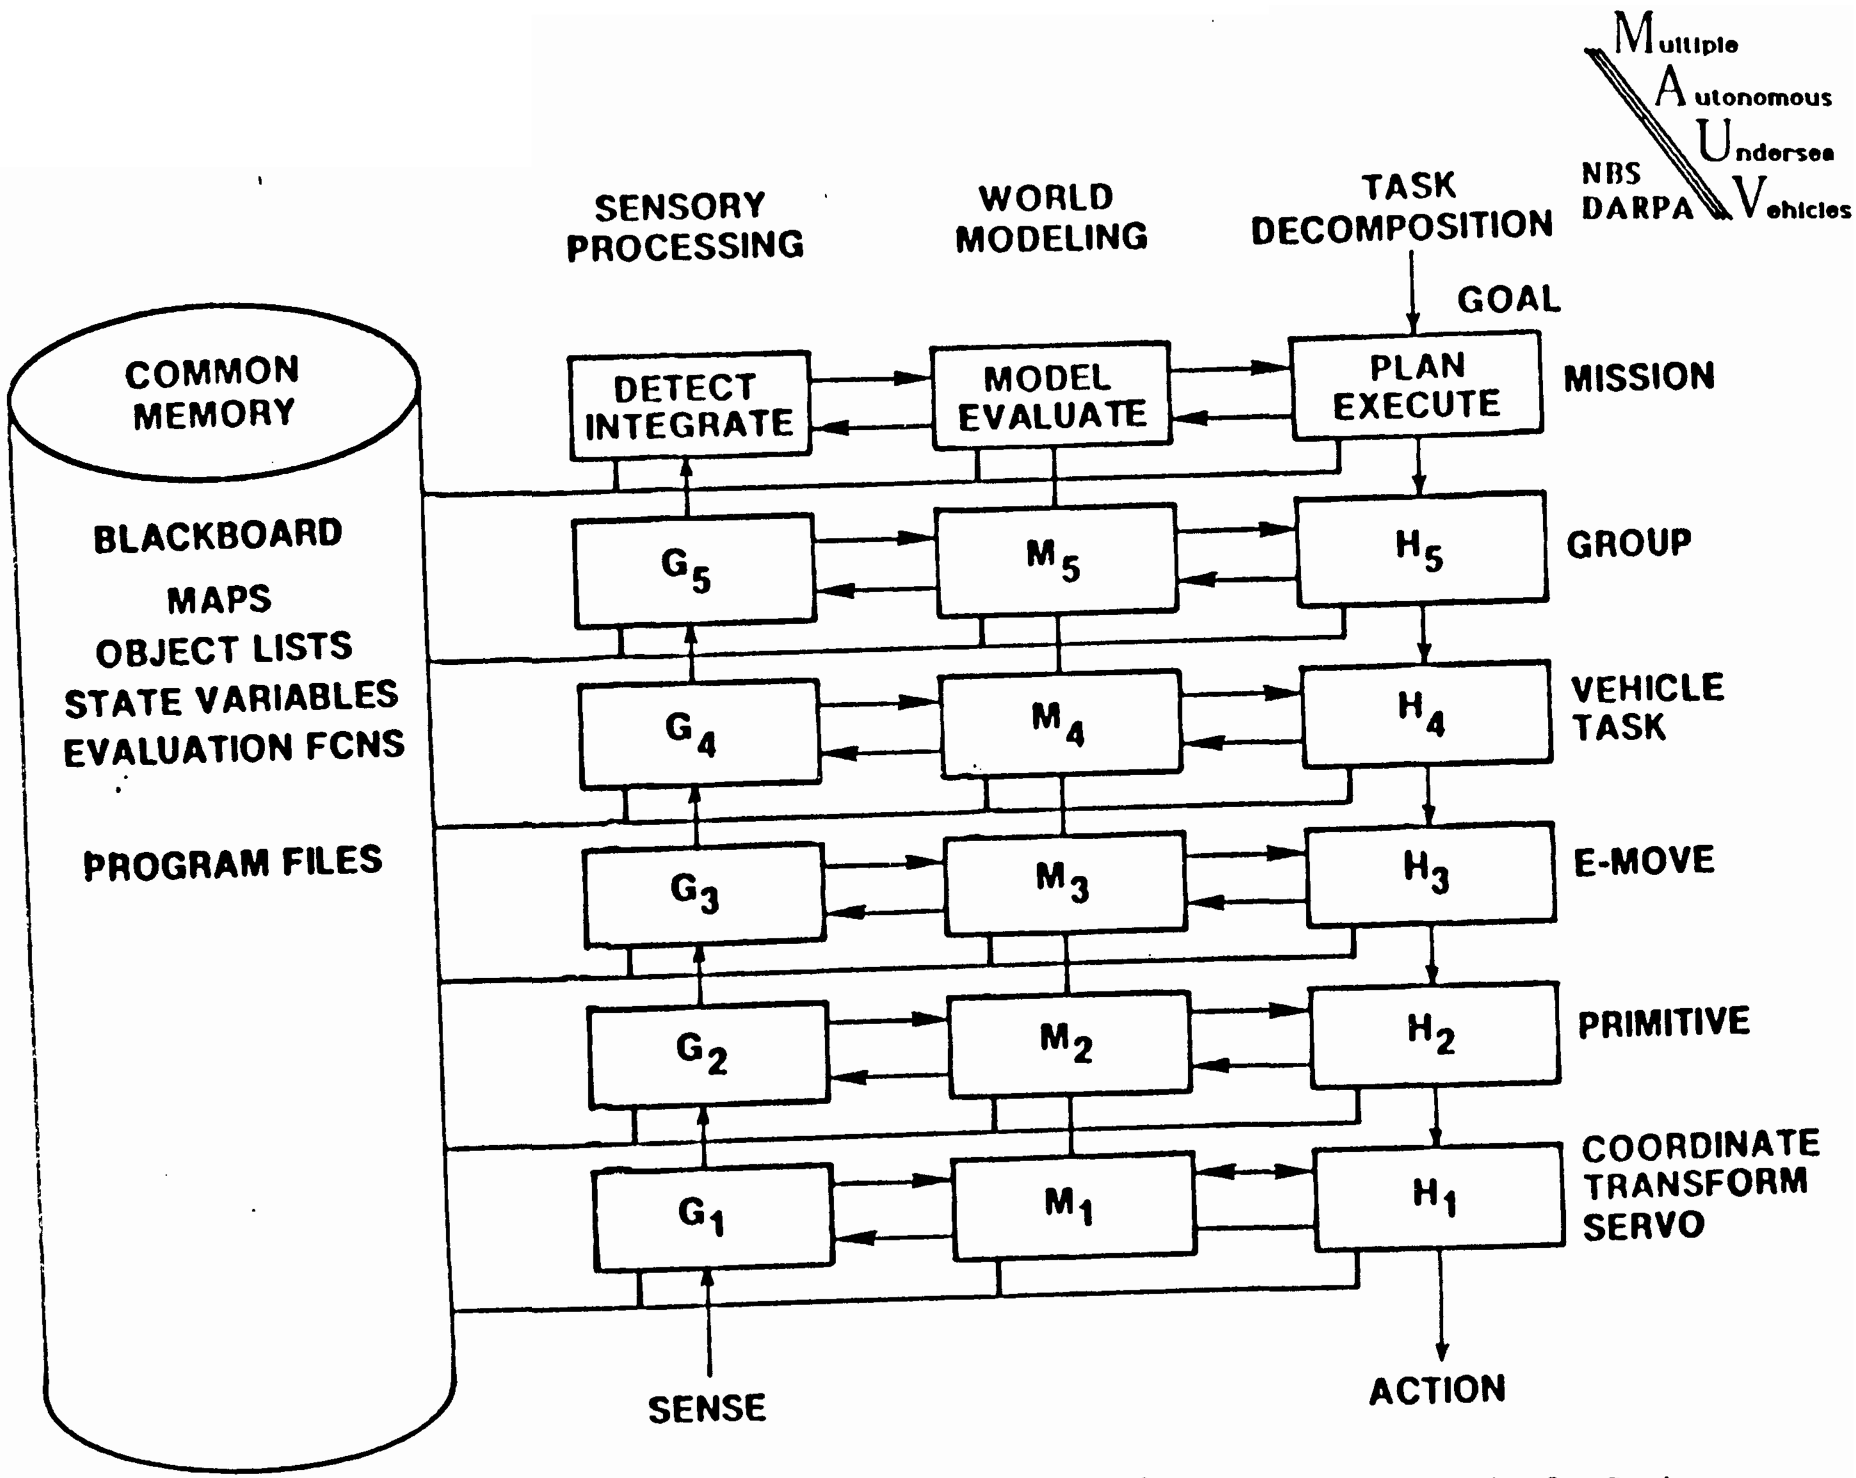

FIGURE 4: A block diagram of the NBS MAUV Control System Architecture.

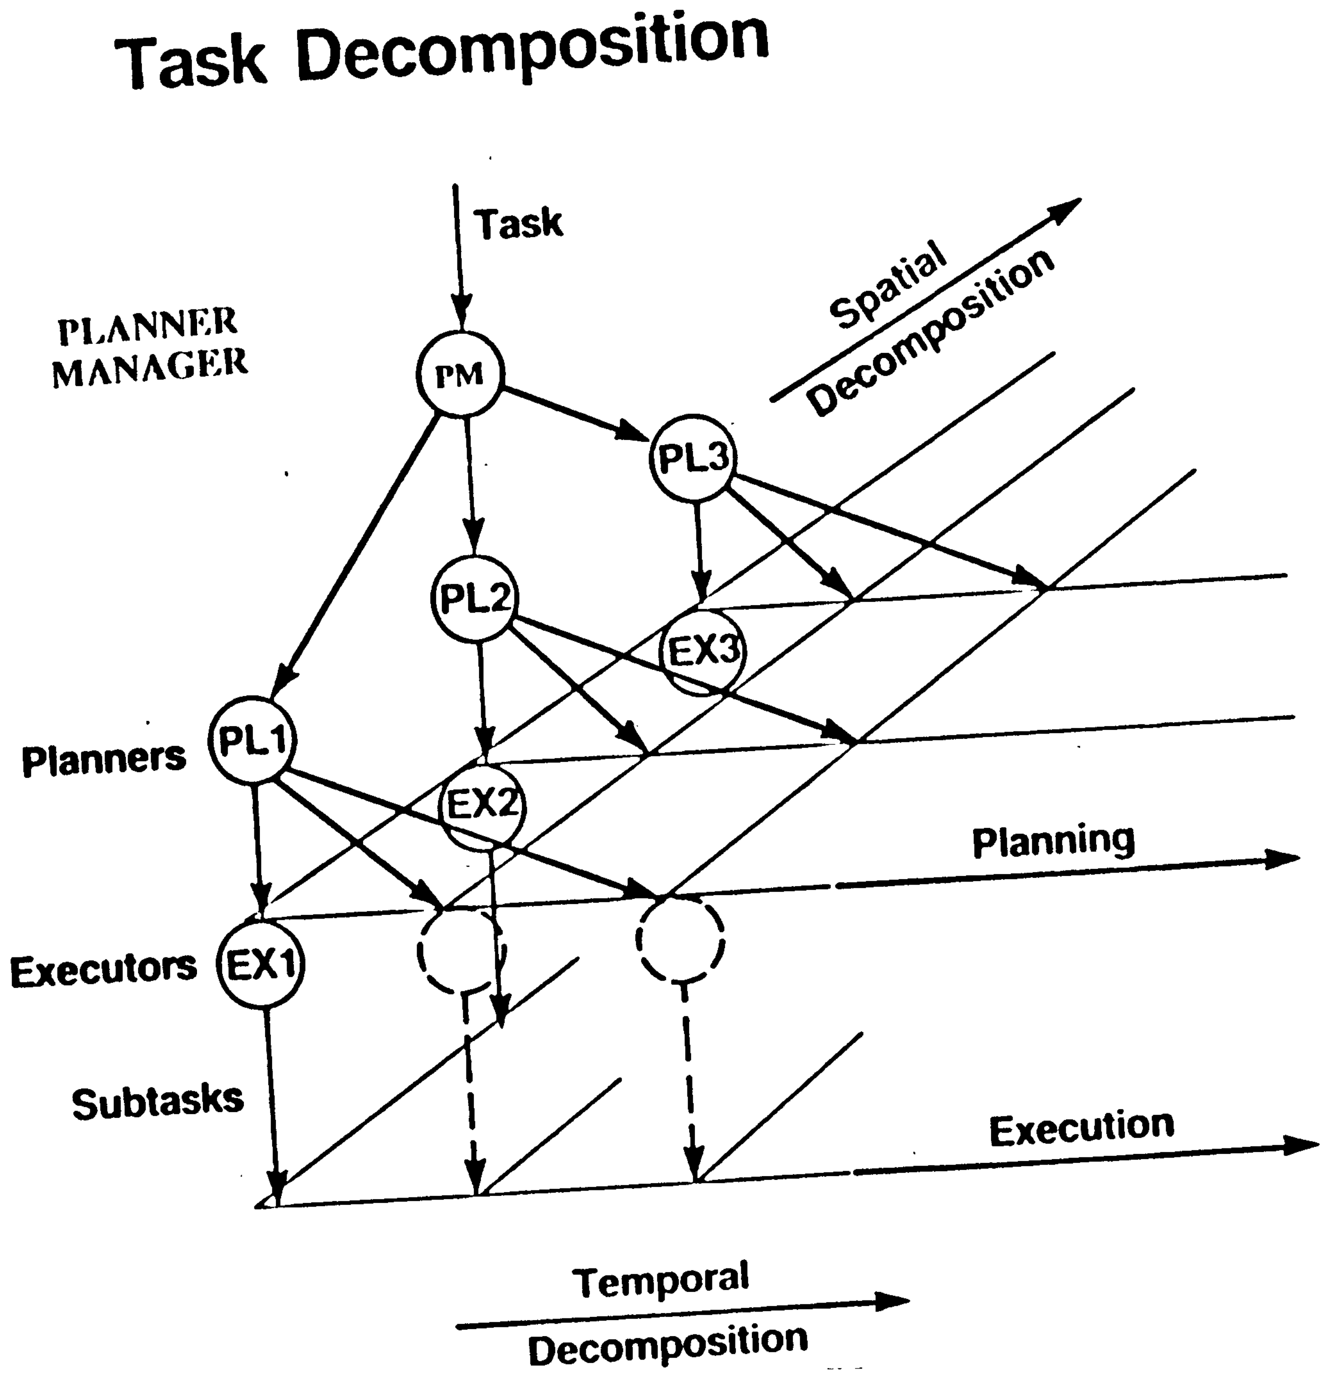

FIGURE 5: Internal structure of the Task Decomposition Modules in the MAUV Control System Architecture at every level of the hierarchy.

World Modeling

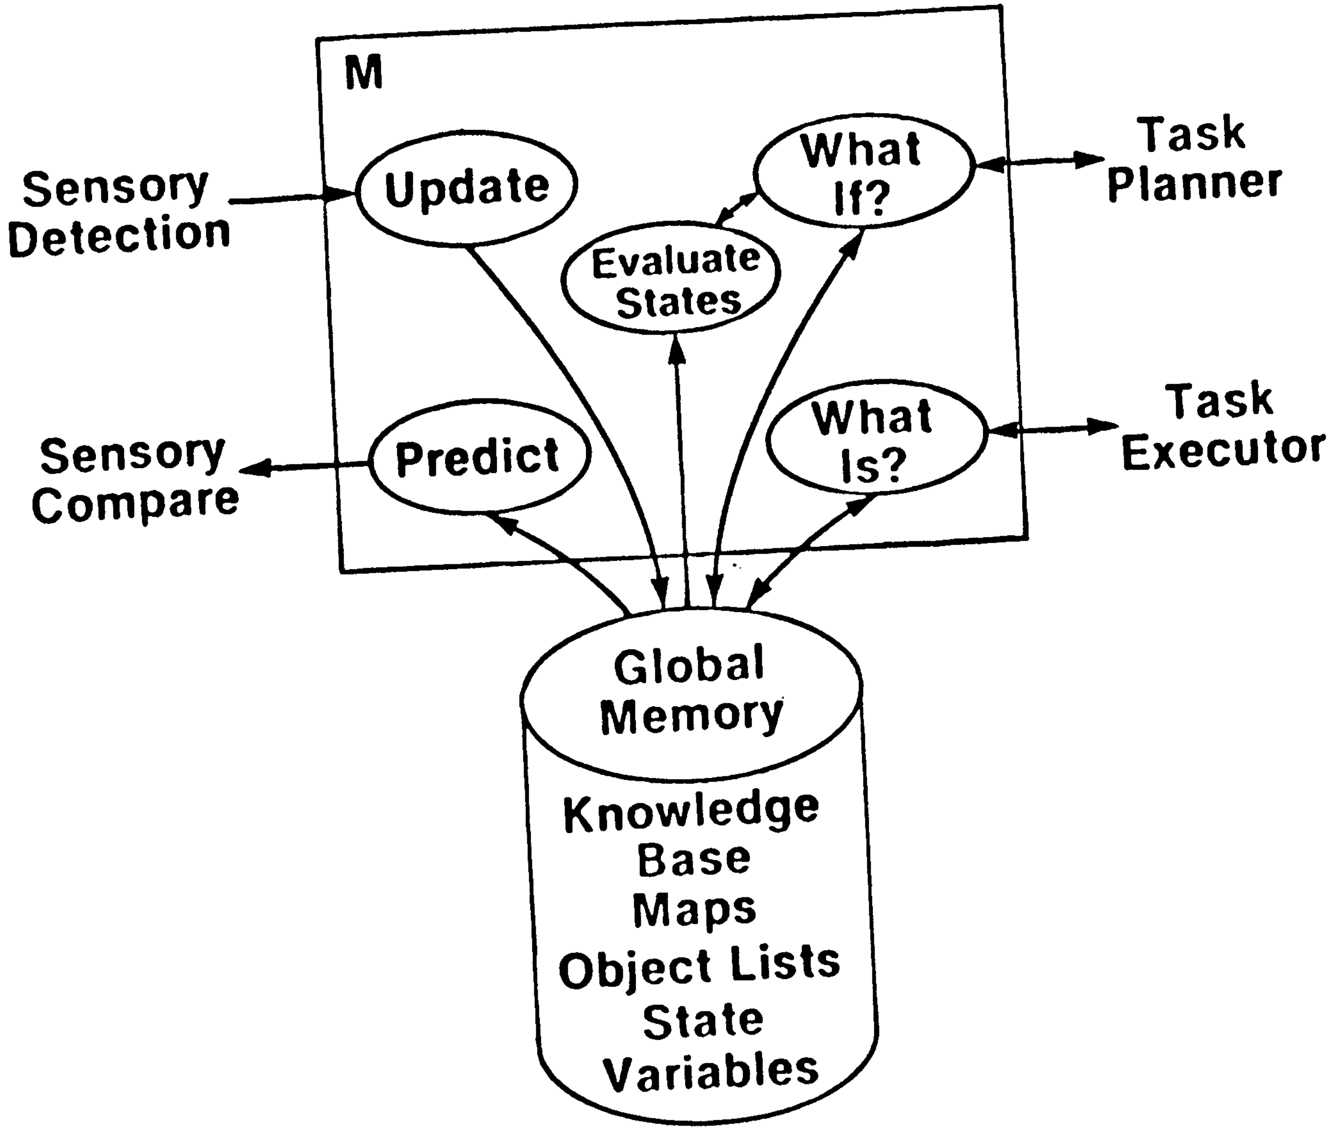

FIGURE 6: Internal structure of the World Modeling Modules at each level.

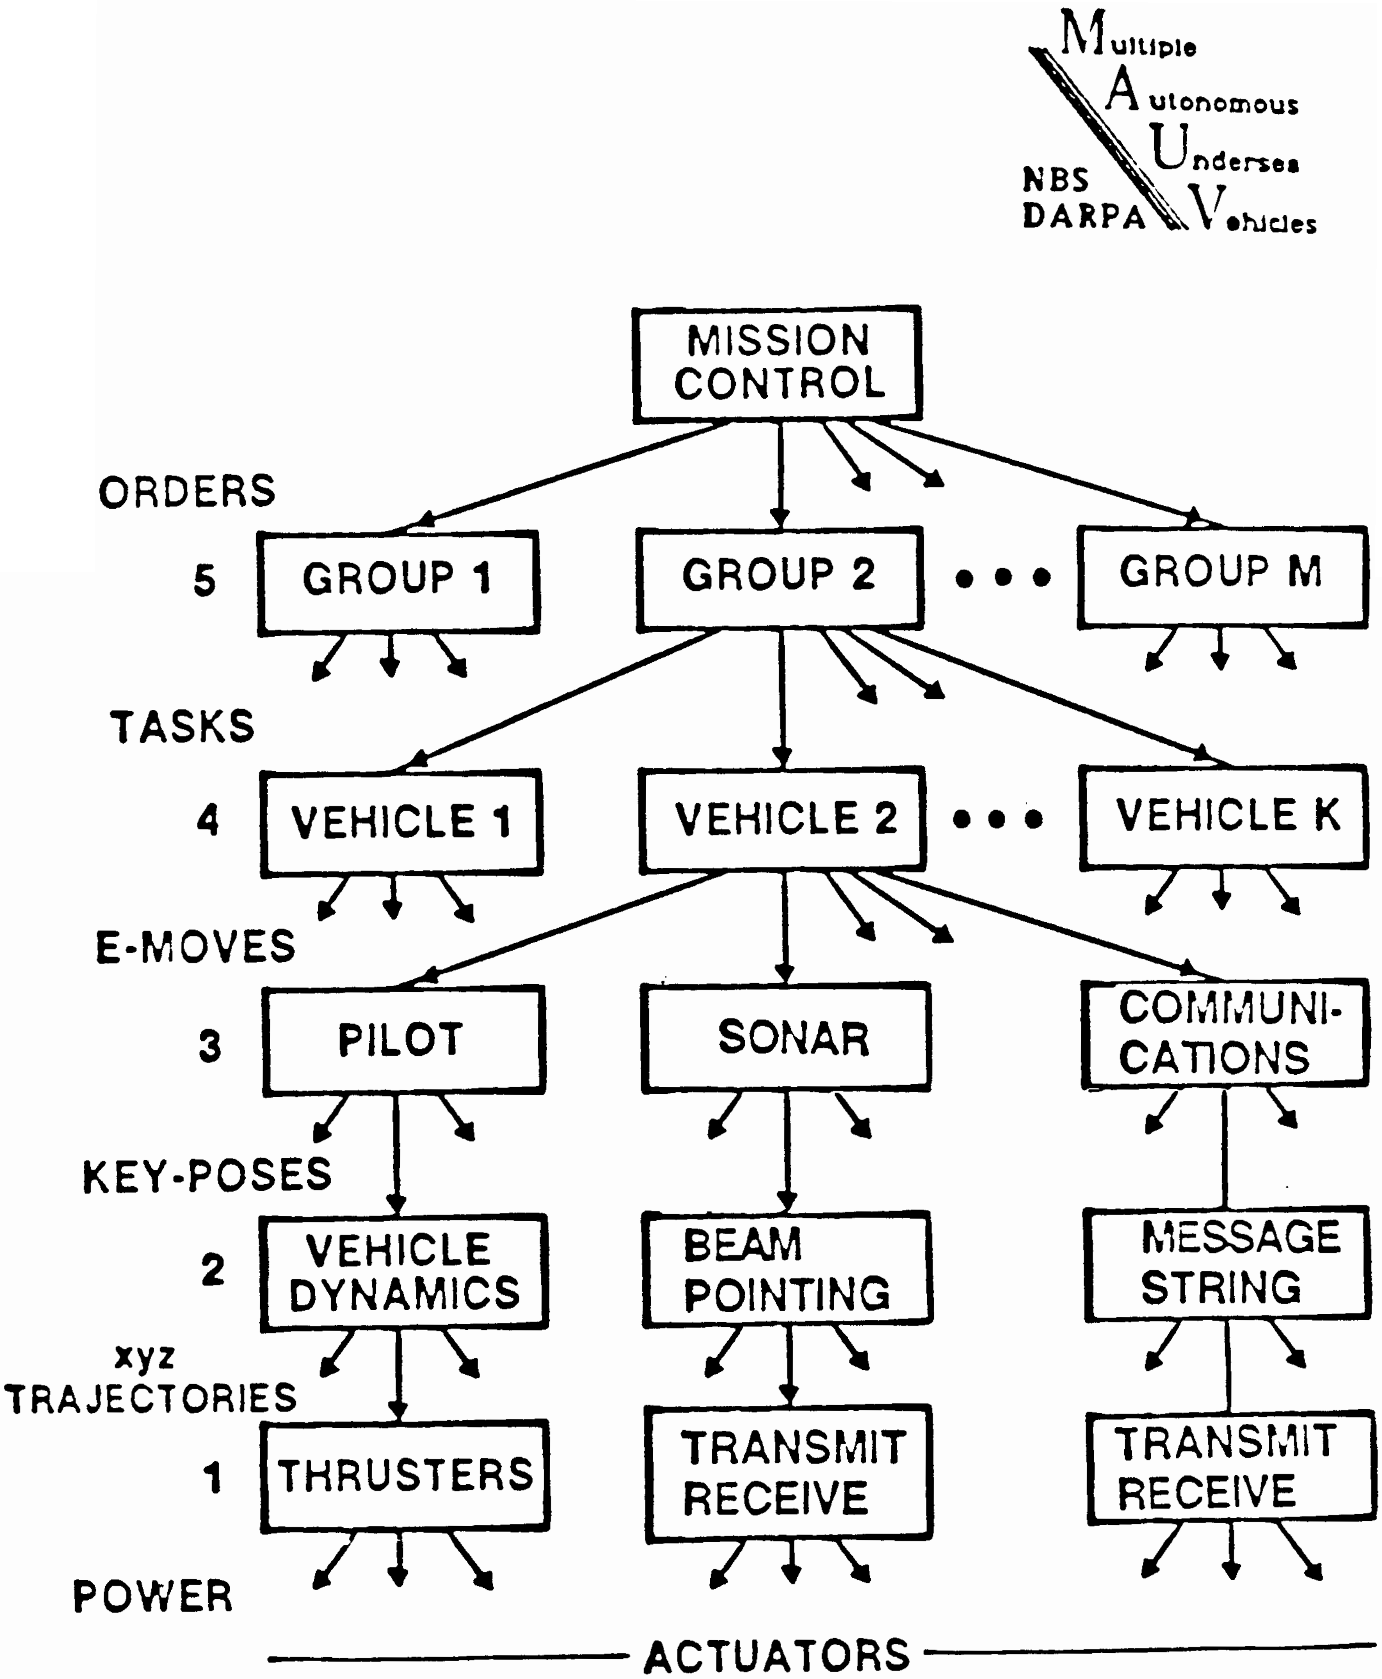

FIGURE 7: Tree structure of the MAUV task decomposition hierarchy (spatial decomposition). The current MAUV project has only one group of two vehicles.

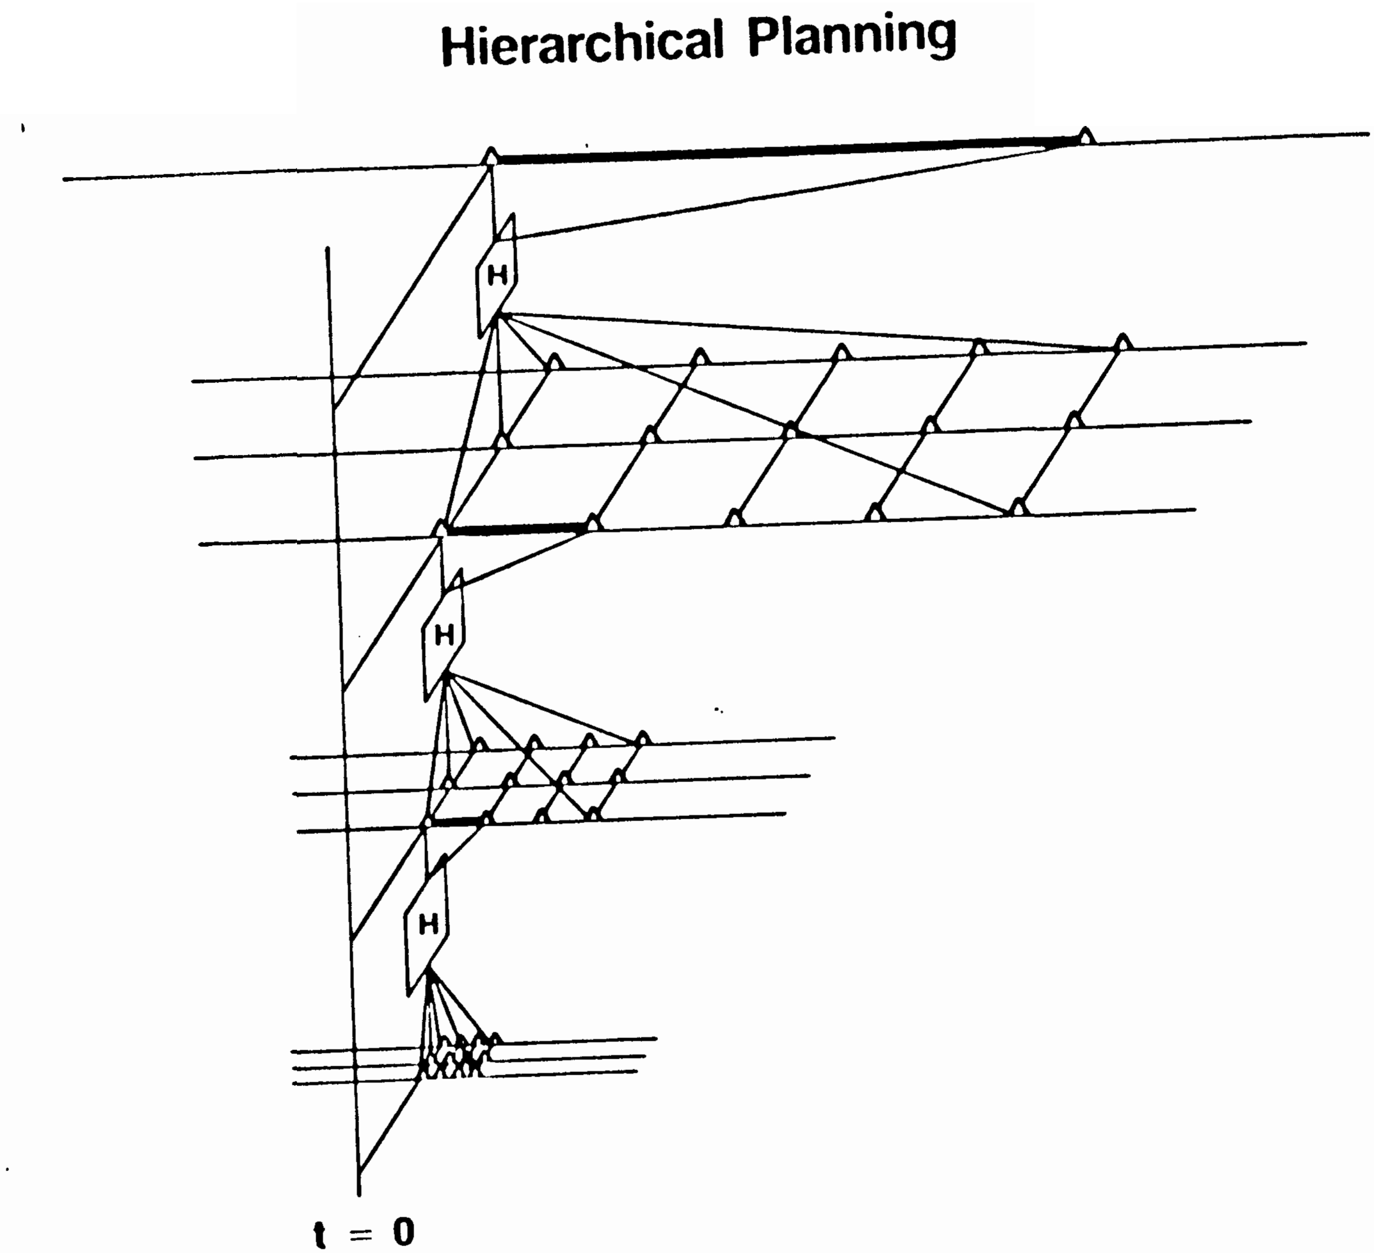

FIGURE 8: Three levels of real-time planning activity in the MAUV hierarchy.

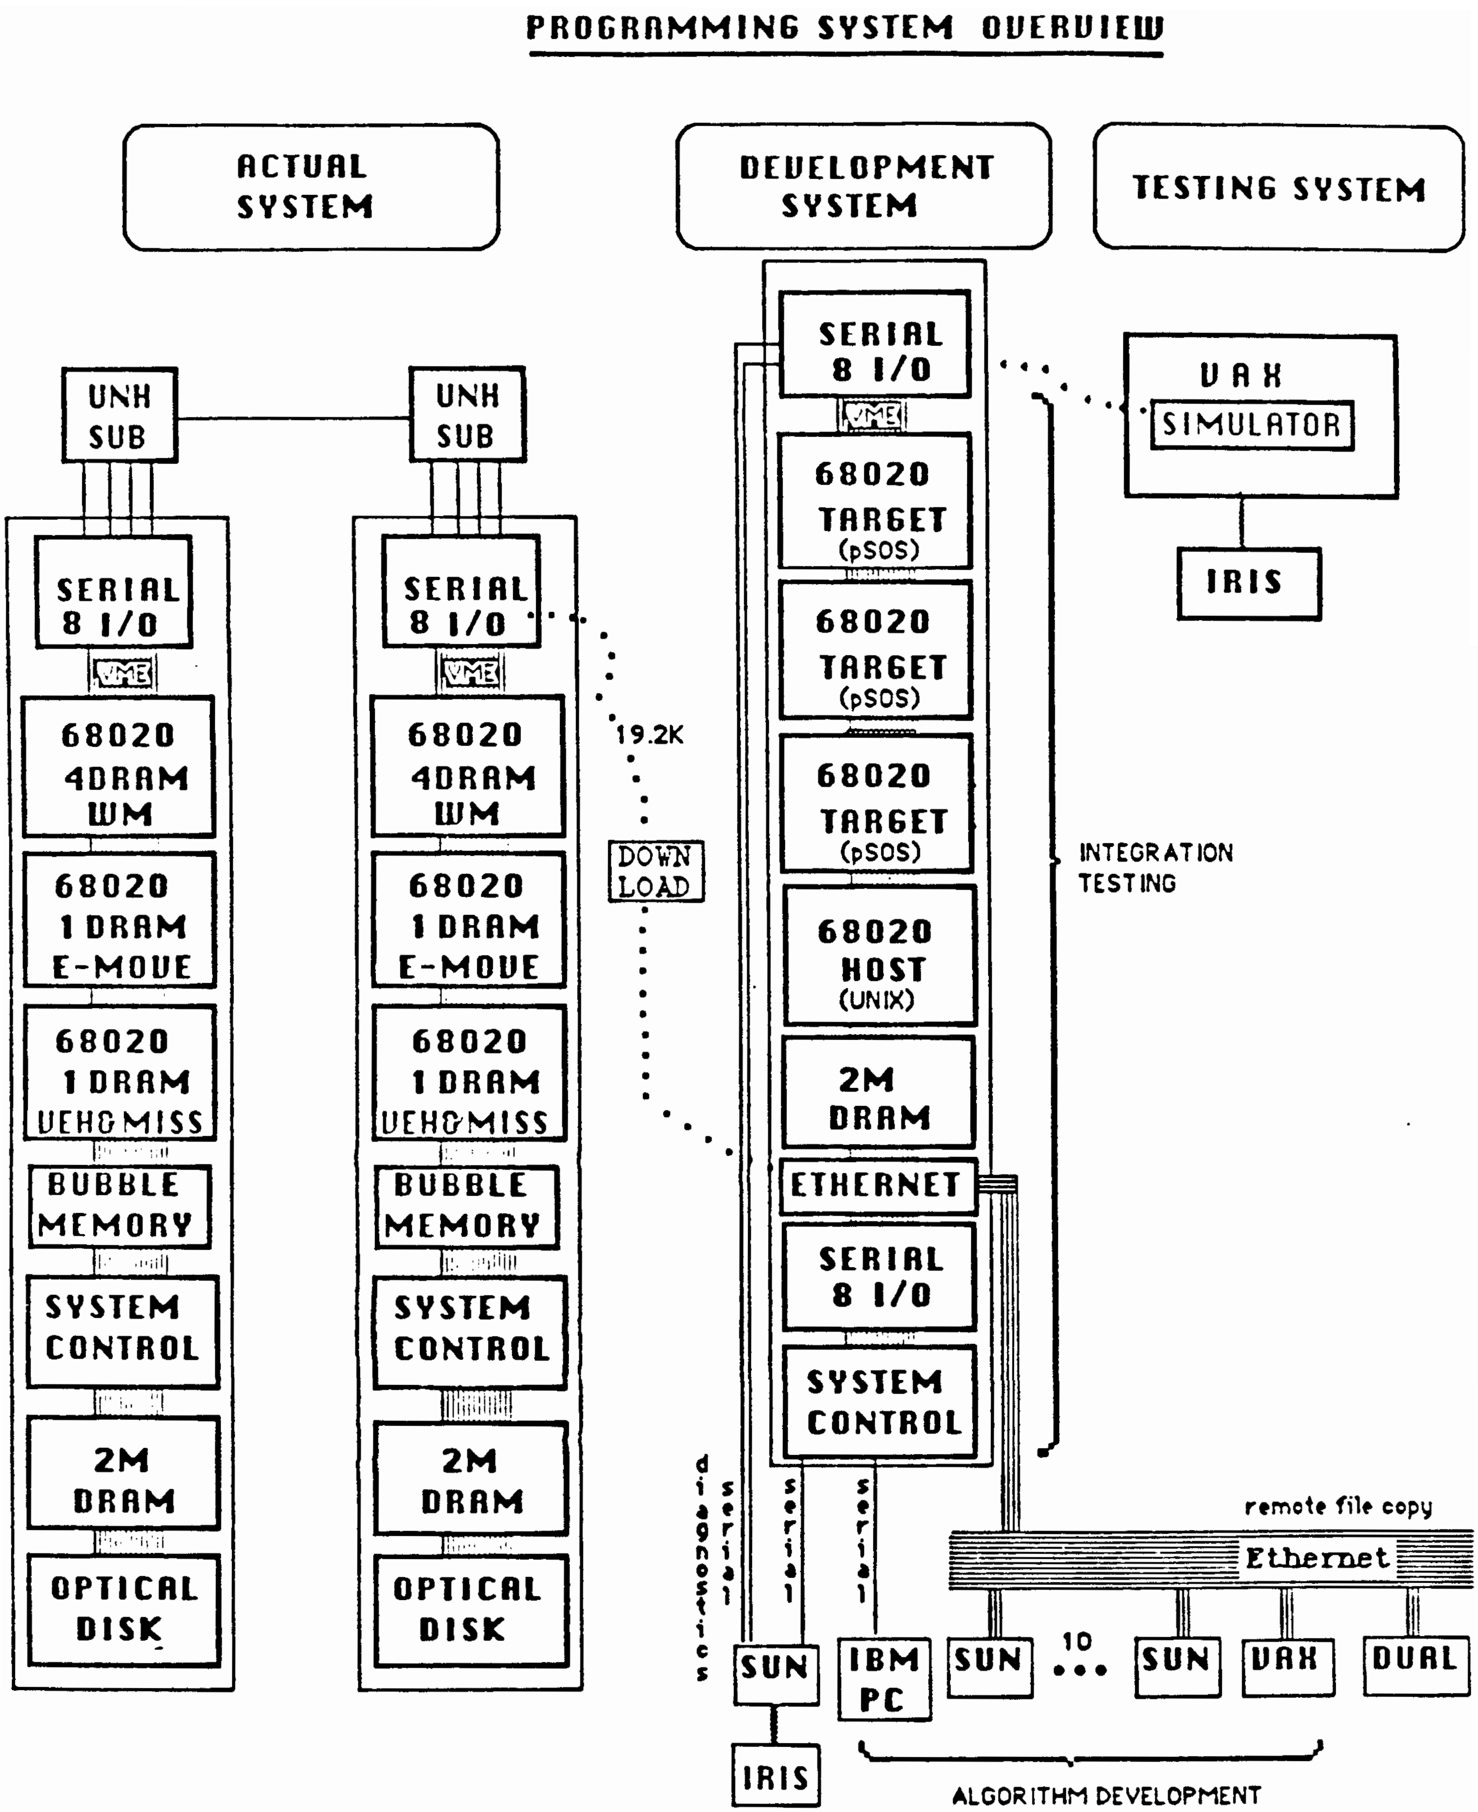

FIGURE 9: On the left is the target hardware for the two MAUV vehicles. On the right is the MAUV software development and simulation environment.